Submitted: Neural Computation

See Globally, Spike Locally: Retinal Oscillations Encode Large Features

Greg J. Stephens[1,3], Sergio Neuenschwander[4], John S. George[1,3], James Theiler[2,3], David W. Marshak[5], Wolf Singer[4], Garrett T. Kenyon[1,3]

[1]Biological and Quantum Physics (P-21), [2]Space and Remote Sensing Sciences (ISR-2), [3]Los Alamos National Laboratory, Los Alamos, NM, 87545; [4]Max-Planck-Institut für Hirnforschung, Deutschordenstraße 46, 60528-Frankfurt am Main, Germany; [5]Department of Neurobiology and Anatomy, University of Texas Medical School, Houston, TX, 77030,

Abbreviated Title: Oscillations Encode Large Features

Corresponding Author:
Garrett T. Kenyon
P-21 Biological and Quantum Physics, MS D454
Los Alamos National Laboratory
Los Alamos, NM 87544
phone: 505-667-1900
fax: 505-665-4507
email: gkenyon@lanl.gov

Acknowledgements: The authors wish to acknowledge useful discussions with Chris Wood and Bryan Travis. This work was supported by the Department of Energy Office of Nonproliferation Research and Engineering, the MIND Institute for Functional Brain Imaging and by the Lab Directed Research and Development Program at the Los Alamos National Laboratory. DWM supported by the National Eye Institute #EY06472 and by the National Institute of Neurological Disease and Stroke #NS38310

Keywords: gamma oscillations, temporal coding, synchrony, computer model

Abstract

Synchronous oscillations evoked by large stimuli have been observed in a variety of vertebrate retinas yet their visual function remains unresolved. Here, we show that coherent oscillations in the spike trains of neighboring retinal ganglion cells encode global topological properties, such as size, that cannot be deduced locally from their mean spike counts or time-averaged firing rates. We analyzed the single-trial responses of cat ganglion cells to spots of various sizes and found that the mean spectral amplitude between 65-100 Hz, computed from multi-unit spike trains 50-300 msec in duration, supports good discrimination (approximately 75%) between large and small spots. These results demonstrate that size information can be extracted from synchronous oscillations on single-trials and short time scales. To examine whether the information content of retinal oscillations depends critically on how such oscillations are generated, as well as to explore an expanded range of stimulus sizes and topologies, we constructed two mathematical models whose *trial-averaged* correlations matched those observed experimentally: 1) a temporally-modulated Poisson process whose mean rate was held constant regardless of stimulus size and 2) an integrate-and-fire feedback circuit based on retinal anatomy. In both models, synchronous oscillations mediated *single-trial* size-discrimination comparable to that measured experimentally. Finally, as retinal oscillations are both size-dependent and stimulus-specific, we expected there to be a sharp transition in the gamma-band activity evoked by random binary images as the pixel density exceeds a critical percolation threshold at which large connected clusters first appear, a novel experimental prediction supported by the retinal model.

Introduction

High frequency oscillations have been observed in a wide variety of vertebrate retinas, including cat (Laufer and Verzeano, 1967; Neuenschwander et al., 1999; Neuenschwander and Singer, 1996; Steinberg, 1966), rabbit (Ariel et al., 1983), mudpuppy (Wachtmeister and Dowling, 1978), frog (Ishikane et al., 1999), macaque (Frishman et al., 2000) and human (De Carli et al., 2001; Wachtmeister, 1998). Nonetheless, the role of synchronous oscillations in normal visual processing is not fully understood. Here, we used experimental and model spike trains to investigate whether coherent oscillations among neighboring retinal ganglion cells encode global stimulus properties not locally available in their mean spike counts or time-averaged firing rates. Ganglion cell firing rates often depend both on local contrast and on stimulus size, increasing for spot diameters smaller than the receptive field center and decreasing thereafter due to surround inhibition. Because it is impossible to disentangle the effects of global stimulus topology from the effects of local contrast, the size of a stimulus cannot be uniquely determined from the mean firing rates of one or a few neighboring cells. High frequency oscillations, on the other hand, are not evoked by small spots (Ariel et al., 1983; Ishikane et al., 1999; Neuenschwander et al., 1999) and thus are not intrinsically confounded by local contrast. Moreover, unlike information encoded by temporal synchrony, which is susceptible to contamination from chance coincidence and therefore potentially unreliable when the firing rate is high (Shadlen and Movshon, 1999), oscillatory activity can be modulated independently of the mean firing rate.

We hypothesize that whereas the average firing rate in a small group of neighboring ganglion cells conveys information about local stimulus properties, such as

contrast, synchronous oscillations convey additional information about global stimulus properties, such as size. To test this hypothesis, we asked whether global topological information could be extracted from the coherent oscillations among neighboring retinal ganglion cells upon single stimulus presentations lasting no more than a few hundred msec, equivalent to typical inter-saccade intervals (Martinez-Conde et al., 2004). The answer to this question both constrains and informs the functional role of retinal oscillations. In order to be behaviorally relevant, any information conveyed by synchronous oscillations must be accessible on physiological time scales. Furthermore, spatial convergence at early visual processing stages is rather low, being one to a few at retina-LGN synapses (Usrey et al., 1999) and approximately one to 15 at LGN synapses onto layer IV simple cells if ON and OFF inputs are considered separately (Alonso et al., 2001). The low convergence onto individual target neurons at early visual processing stages requires that any information encoded by synchronous oscillations must be obtainable from a relatively small number of cells. Conversely, if size-dependent oscillations cannot be extracted from a relatively small number of ganglion cells on single stimulus trials, then it is unlikely that such oscillations could be functionally relevant at early processing stages.

To investigate whether global stimulus properties could be extracted from the synchronous oscillations between a relatively small number of retinal ganglion cells on single-trials, we re-analyses multi-unit spike trains from the cat retina recording in response to spots of different sizes. A discrimination analysis was used to classify the relative sizes of spot stimuli as either “smaller” or “larger” on the basis of oscillations present in short sections of multi-unit spike train data, from 50 to 300 msec in duration

and containing on the order of 4 to 5 ganglion cells located at the center of the stimulus. Frequency spectra were computed from the individual epochs of multi-unit data and the average spectral amplitude in the upper gamma-band, between 65 and 100 Hz, was measured relative to baseline after normalizing by the total number of spikes. Despite the presence of substantial single-trial variability, it was possible to discriminate between large and small stimuli with performance levels of approximately 75%.

An additional question we sought to address was whether the single trial discrimination mediated by experimentally recorded spike trains resulted from highly specialized processing or was instead a general property of coherent oscillations that would likely be present regardless of the underlying generating mechanism. Simply because two processes exhibit similar multi-trial correlations does not require that their single-trial statistics will likewise be identical. We therefore analyzed an equivalent single-trial discrimination task using two very different models for generating synchronous oscillations.

Model I: Rate-Modulated Poisson Process (Poisson). Temporal correlations between the separate elements contributing to the multi-unit spike train were generated by a common oscillatory modulation of their instantaneous firing probabilities although the elements themselves did not interact. The average number of spikes produced by the Poisson model was held constant across stimulus conditions to ensure that no size information was conveyed by the mean firing rate.

Model II: Retinal Feedback Circuit (Network). Coherent oscillations were generated by feedback interactions between ganglion cells and amacrine cells in a manner consistent with retinal anatomy and physiology (Kenyon et al., 2003).

Both models were calibrated to produce size-dependent synchronous oscillations that were similar to those measured experimentally, as assessed by *trial-averaged* multi-unit cross- and auto-correlation histograms as well as by multi-unit frequency spectra. Despite the very different mechanisms used to generate synchronous oscillations in the two sets of artificial spike trains, both models mediated good single-trial size discrimination consistent with that inferred from the experimental data.

Artificially generated spike trains were also used to investigate the amount of potentially useful size information conveyed by synchronous oscillations among varying numbers of neurons, from 1 to 16, and to explore the range of stimulus sizes that could be discriminated based on gamma-band activity alone. By employing two separate models based on very different mechanisms for generating synchronous oscillations, we were able to draw stronger inferences than would be possible with only a single model. Lastly, the network model was used to examine oscillatory activity across the entire retina in response to a set of random binary images. Each image was characterized by an average pixel density, or probability that any given pixel was on. Consistent with the encoding of contiguous size by retinal oscillations, there was a sharp transition in the total gamma-band activity near the image percolation threshold at which large connected clusters first appear. The observation of a sharp transition in the total gamma-band activity near the percolation threshold represents a novel experimental prediction.

Methods

Physiological Recordings:

Data from previously reported experiments in the cat retina were reanalyzed in order to quantify the information relating to the global stimulus topology conveyed by high frequency oscillations in short sections of multi-unit spike train data. As no new experiments were conducted, and since a detailed description of the experimental methods is available elsewhere (Neuenschwander et al., 1999), only an abbreviated description of the experimental procedure is provided here.

Intra-ocular recordings were made from anesthetized and paralyzed cats in response to spots of various sizes, presented at high contrast against a background illumination of ~0.4 cd-m^{-2}. Each spot of a given size was presented 20 times for a period of three seconds. For each spot size, the total spike record across all trials was divided into 200 equally spaced, non-overlapping segments, from 50-300 msec in length. For purposes of this study, these 200 spike train segments were treated as independent stimulus trials.

Model I: Rate-Modulated Poisson Process (Poisson):

Artificial spike trains were generated under the assumption that the spatiotemporal correlations between cat ganglion cells are due to a common oscillatory input that uniformly modulates the firing rates of all simultaneously recorded neurons. An oscillatory time series of a duration, T, and temporal resolution, Δt, could be constructed by first defining the discrete frequencies, f_k:

$$f_k = \frac{k}{T}, 0 \le k < \frac{T}{\Delta t} \quad (1)$$

in terms of which the discrete Fourier coefficients were defined as follows:

$$C_k = C\exp(2\pi i r_k)\exp\left(\frac{(f_k - f_0)^2}{2\sigma^2}\right) \quad (2)$$

where f_0 is the central oscillation frequency, σ is the width of the spectral peak in the associated power spectrum, r_k is a uniform random deviate between 0 and 1 that randomized the phases of the individual Fourier components (generated by the Matlab® intrinsic function RAND) and C is an overall scale factor. The coefficients, C_k, were used to convert back to the time domain using the discrete inverse Fourier transform:

$$R_n = A\frac{1}{N}\sum_{k=1}^{N-1} C_k e^{-2\pi i f_k t_n} + R_0 \quad (3)$$

where the real part of R_n denotes the value of the time-dependent firing rate at the discrete times, $t_n = n\,\Delta t$, $N = T\,\Delta t$, A is an empirically determined scale factor and we have added a constant offset, R_0, which sets the mean firing rate, equal to 50 Hz. The quantity A, with units of Hz, was determined by the formula: $A = R_0\,(0.16 + 0.14L)$, where L represents the linear stimulus size, or length, along one axis, with values ranging from 1 to 16, and the other coefficients were determined empirically to produce a reasonable match to the experimental data. Stimulus sizes denoted "small", "med" and "large" in the text corresponded to L values of 1, 4 and 6, respectively. The quantity C in eq. 2 was determined by setting the standard deviation of R_n over all time steps to unity when $A = 1$. The width of the frequency spectrum, σ, also in units of Hz, was given by an analogous formula: $\sigma = 9.4 - 0.6L$. Negative values of R_n were truncated at zero and the resulting time series rescaled so that its average value remained equal to R_0.

The time series defined by R_n was used to generate oscillatory spike trains via a pseudo-random process:

$$S_n = \theta\left(R_n \Delta t - r\right) \quad (4)$$

where $R_n \Delta t$ is the probability of a spike in the nth time bin, θ is a step function, $\theta(x<0) = 0$, $\theta(x \geq 0) = 1$, and r is again a uniform random deviate. In the limit that $R_n \Delta t << 1$, the above procedure reduces to a rate-modulated Poisson process. The same time series, R_n, was used to modulate the firing rate of each element contributing to the artificially generated multi-unit spike train, thus producing temporal correlations due to co-varying input.

Model II: Integrate-and-Fire Feedback Circuit (Network):

A second and very different method for generating artificial spike trains used a semi-realistic model of the inner retina (Kenyon et al., 2003). Input to the model was conveyed by an array of external currents proportional to the pixel-by-pixel grayscale value of a two-dimensional image. These external currents directly stimulated the model bipolar cells and approximated their light-modulated synaptic input from cone photoreceptors. The bipolar cells produced excitatory postsynaptic potentials in both ganglion cells and amacrine cells according to a random process (Freed, 2000). The axon-bearing amacrine cells were electrically coupled to neighboring ganglion cells and to each other, and made strong inhibitory connections onto the surrounding ganglion cells and axon-bearing amacrine cells (fig. 1). This feedback circuit produced robust, physiologically realistic oscillations in response to large stimuli. When several ganglion cells were activated by a stimulus, they in turn activated neighboring axon-bearing amacrine cells via gap junctions (Dacey and Brace, 1992; Jacoby et al., 1996; Vaney, 1994). The stimulated cells were then hyperpolarized by the ensuing wave of axon-

mediated inhibition, thus setting up the next cycle of the oscillation. Spike generation was modeled as a leaky integrate-and-fire process with a membrane time constant of 5 msec, consistent with published physiological data from cat alpha ganglion cells (O'Brien et al., 2002). The model also contained local non-spiking amacrine cells that generated randomly distributed inhibitory postsynaptic potentials that helped make spontaneous firing asynchronous, but were not otherwise critical for generating oscillatory spike trains. Mathematical details of the model are presented in the appendix. A thorough analysis of the behavior of the model, its robustness with respect to both numerical and physiological parameters and its connection to experimental data, is presented elsewhere (Kenyon et al., 2003).

Figure 1 about here

Correlation Analysis

Using multi-unit spike trains recorded from the cat retina, multi-unit auto-correlograms were computed for each segment, 200 msec in duration, and the individual correlation functions normalized as a fraction of the expected synchrony due to chance, defined as the mean correlation amplitude for delays between 75-100 msec. A correlation amplitude of one thus corresponded to a doubling in the number of correlated events over the expected rate during that segment. By measuring correlations relative to the baseline present during the same trial, it was possible to compensate, in part, for changes in firing activity during the 3 second stimulus presentation. All correlations were lag corrected for edge effects arising from the finite length of the spike train. The single-trial auto-correlograms were averaged to produce a final correlation measure.

A similar analysis was used to measure correlations in artificially generated spike train segments, also 200 msec in duration. However, for the artificially

generated segments, correlations were computed separately between all distinct cell pairs contributing to the multi-unit record and the individual cross-correlations then averaged to produce a multi-unit cross-correlogram that, unlike the multi-unit auto-correlograms computed from the experimental data, remained finite at zero delay. The multi-unit cross-correlograms measured on each stimulus trial were normalized relative to the expected correlations due to chance and averaged across all stimulus trials. All multi-trial correlation measures used a bin width of 1 msec. Shift-predictors were not subtracted, as our hypothesis does not depend on whether oscillations are phase-locked to the stimulus onset.

Size Discrimination

We assessed the ability to discriminate between small and large stimuli based on the average spectral amplitude in the upper gamma frequency band (65-100 Hz) as extracted from short segments of multi-unit spike train data on single trials. For multi-unit spike trains recorded from the cat retina, individual segments were 50-300 msec in duration and contained approximately 4-5 neurons, the exact number being unknown. Artificially generated spike train segments were 50-400 msec in duration and contained from 1-16 neurons, as indicated. For spike trains generated by the integrate-and-fire retinal model, segments were always extracted from separate stimulus presentations during the plateau portion of the response, beginning 200 msec after stimulus onset. To estimate the average spectral amplitude in the upper gamma band, the multi-unit spike train was Fourier transformed and the mean amplitude of the discrete Fourier coefficients between 65 and 100 Hz was computed. The individual Fourier coefficients were normalized by the total power in the zero-frequency band, thus removing any direct dependence on the mean firing rate. The average spectral amplitude in the upper gamma band was expressed as a fraction of the mean baseline spectral amplitude, defined here as

the average between 220 and 500 Hz, computed on the same trial. This normalization was not essential for the artificially generated spike trains but was necessary to compensate for the changing baseline of the experimentally recorded data. The average spectral amplitude in the upper gamma band (gamma-band activity) was determined for each trial and the results sorted into 11 equally spaced bins spanning the full range of the data. The binned, single-trial gamma-band activity was then normalized as a probability distribution, yielding the fraction of trials on which the single-trial gamma-band activity fell within a particular range.

Given the probability distributions of single-trial gamma-band activity for each stimulus size, it was straightforward to estimate the percentage of trials on which any given pair of stimulus sizes could be discriminated based on the single-trial gamma-band activity alone. Specifically, the ability to discriminate two stimulus sizes was inversely related to the degree of overlap between the corresponding probability distributions (Duda et al., 2001). If the distributions of single-trial gamma-band activity overlapped completely, the maximal theoretical performance on the size discrimination task would be no better than chance (50% correct). On the other hand, if the distributions of single-trial gamma-band activity were entirely non-overlapping for a given pair of stimulus sizes, the maximum theoretical performance on the discrimination task would be perfect (100% correct). Between these two extremes, corresponding to distributions that partially overlap, maximum theoretical performance on the size discrimination task, P, expressed as a fraction of trials correctly classified, is given by the following formula:

$$P = \left(2 - A_{overlap}\right)/2 \tag{5}$$

where $A_{overlap}$ denotes the total area of the overlap between the two distributions and the maximum value of $A_{overlap}$ is normalized to one. Error bars on the estimated values of P

were determined by assuming the number of trials to either side of the Bayes discriminator obeyed binomial statistics.

Results

As reported previously (Neuenschwander et al., 1999), multi-unit auto-correlograms recorded from the cat retina, computed from multi-unit spike trains containing approximately 4-5 ganglion cells, exhibit synchronous high frequency oscillations whose amplitude increases markedly with stimulus size (fig. 2a). A small square spot, 0.7°×0.7°, produced negligible oscillations, whereas a large square spot, 9.8°×9.8°, evoked strong oscillations at approximately 80 Hz. An intermediate sized spot, 6.3°×6.3°, produced periodic correlations at similar frequencies to those elicited by the larger spot but of smaller amplitude and duration, indicating that over this range of spot dimensions oscillatory responses depend smoothly on stimulus size. Periodic modulations in the multi-unit auto-correlogram declined with increasing delay, indicating that the oscillation phase drifts randomly over time. The multi-unit auto-correlogram is undefined at zero delay and was set arbitrarily to zero.

Figure 2 about here

Multi-unit cross-correlograms computed from artificial spike trains produced by either by a rate-modulated Poisson process (model I, Poisson, fig. 2b) or by a semi-realistic retinal feedback circuit (model II, Network, fig. 2c) exhibited oscillations whose size dependence, amplitude, duration and frequency were similar to those measured experimentally. For the rate-modulated Poisson process, the amplitude and width of the spectral peak were adjusted to produce a good qualitative fit to the physiological data for each spot size, labeled “small”, “med” and “large”. For the integrate-and-fire feedback circuit, we used square spots covering either 1×1, 4×4, or 6×6 ganglion cells (stimulus

intensity = 0.25). Poisson distributed events were added to the spike trains generated by the network model so as to increase the mean firing rate by 20% regardless of stimulus size. This procedure yielded correlations that better matched the experimental data, possibly because multi-unit spike trains recorded from the cat retina may have contained some non-oscillatory background spikes. Both of the artificially generated event trains consisted of 4 units, which in the case of the retinal model were taken from a 2×2 array of neighboring ganglion cells located at the center of the stimulus.

To quantify the dependence of oscillatory responses on stimulus size, we Fourier analyzed 200 msec epochs of both natural and artificial multi-unit spike trains, using the same data as for the correlation analysis above. The magnitudes of the Fourier coefficients computed from individual spike train segments were normalized by the total number of events, given by the Fourier coefficient at 0 Hz, and the results from all 200 stimulus trials averaged together (fig. 3). In all three sets of data, a prominent peak in the upper gamma frequency band, defined here as between 65-100 Hz, became evident as the size of the stimulus increased. The frequency spectra associated with the rate-modulated Poisson process (model I) were somewhat broader than those computed from either the experimental data or from the retinal network (model II), even though the multi-unit correlograms derived from all three sources were very similar (fig. 2). Differences between the frequency spectra produced by the two models suggest that the retinal network (model II) better describes for the physiological mechanisms underlying synchronous oscillations.

Figure 3 about here

The measures of oscillatory responses among retinal ganglion cells presented above, either multi-unit correlograms or frequency spectra, represent averages over multiple trials. On behavioral time scales, however, such multi-trial measures are not available to the animal's nervous system. For retinal oscillations to convey behaviorally relevant information, they must be detectable on single trials lasting from 10s to 100s of msecs. For both the experimental and artificially generated spike trains, we examined the normalized frequency spectra computed from single stimulus trials, 200 msec in duration, in response to three different spot sizes (fig. 4). As in the multi-trial averages, spectral amplitudes between 65-100 Hz measured from single-trial data increased systematically as a function of spot size. In the cat data, the single-trial frequency spectra were substantially noisier in response to the smallest spot size (0.7°×0.7°), presumably because the receptive field centers were not completely covered, thereby reducing the overall number of spikes compared to the larger spot sizes.

Figure 4 about here

To investigate the size information encoded by retinal oscillations on behavioral time scales, we constructed probability distributions of the gamma-band activity present on single-stimulus trials (fig. 5). The gamma-band activity was defined as the average spectral amplitude between 65-100 Hz, normalized by the DC component and expressed as a fraction of the average spectral amplitude between 220-500 Hz present during the same 200 msec epoch. The analysis was repeated for all 200 trials and the results sorted into 11 uniform bins spanning the full range of the data. For both natural and artificial spike trains, the distributions for the smallest (solid lines) and largest (dotted lines) spot sizes were reasonably distinct, implying that gamma-band activity from a small group of

neighboring ganglion cells can mediate good single-trial size discrimination. For the cat spike trains, the distribution of gamma-band activity evoked by the smallest spot size was shifted somewhat to the right compared to the corresponding distributions from the two models, reflecting the higher level of noise present in the single-trial spectra computed from the experimental data (fig. 4a). Discrimination between small and large spots is thus expected to be slightly worse for the experimental vs. model spike trains.

Figure 5 about here

To quantify the degree of discrimination theoretically possible on single stimulus trials, we employed a Bayes discriminator (Duda et al., 2001). For any pair of spot sizes, the Bayes discriminator is defined as the value of the gamma-band activity at which the two distributions cross (this definition does not apply if the distributions have multiple intercepts). If the single-trial gamma-band activity falls to the left of the Bayes discriminator, the stimulus is classified as the smaller of the two spot sizes, whereas the stimulus is classified as the larger of the two if the gamma-band activity falls to the right of the discriminator. The overlap between the two distributions estimates the number of trials that, on average, will be classified incorrectly. Using the probability distributions described above, we found that 200 msec segments of multi-unit spike trains from the cat retina permit discrimination between large and small spots on approximately 75% of the trials (fig. 6). Thus, the gamma-band activity from a small number of neighboring ganglion cells conveyed enough information on behaviorally relevant time scales to determine whether the recorded neurons were responding to the “smaller” or to the “larger” stimulus. The artificially generated spike trains yielded slightly better performance on the size discrimination task, exceeding 80% correct for both models

when deciding between the large and small spots. The spike trains from model I, corresponding to a rate-modulated Poisson process, all had the same mean firing rate regardless of stimulus size, demonstrating that differences in the total number of spikes were not essential for the size discrimination task.

Figure 6 about here

It is uncertain why the spike trains generated artificially yielded better performance then those recorded experimentally. One explanation is that the experimentally recorded spike trains may have included ganglion cells of different types and opposite ON/OFF response polarities whereas the computer generated spike trains were derived from a homogeneous population. The changing baseline of the experimentally recorded spikes trains may also have confounded the discrimination analysis, despite efforts to correct for non-stationary effects (see methods). Thus, the performance levels obtained with the experimentally recorded data likely reflect a lower bound on the maximum performance levels theoretically attainable.

We next examined how the information about the stimulus size encoded by retinal oscillations depends both on the duration of, and number of neurons contributing to, the multi-unit spike record. Performance on the size discrimination task, as measured by the percentage of correctly classified trials in which either the small or the large spot was presented with equal probability, was plotted as a function of the analysis window size. For the artificially generated spike trains, the analysis used either 1×1, 2×2, 3×3, or 4×4 neurons, corresponding to 1, 4, 9, or 16 total cells, respectively (fig. 7). For simultaneously recorded cat ganglion cells, performance ranged from approximately 70-75% as the analysis window size increased from 50 to 300 msec. Both theoretical

models suggested that the information content of simultaneously recorded spike trains depends strongly on the number of cells included in the multi-unit record. Even for analysis windows as short as 50 msec, performance levels of up to 85% could be achieved by including a sufficient number of cells. The ability to extract topological information from 50 msec spike train segments was particularly evident for the rate-modulated Poisson process, whereas artificial spike trains generated by the retinal model conveyed substantially less size information on short time scales. This discrepancy may be related to the synchronization of neighboring ganglion cells in the network model due to extensive gap junction coupling via amacrine cells (Kenyon et al., 2003), since such synchronization effectively reduces the improvement in signal to noise obtained by pooling over multiple neurons (Shadlen and Newsome, 1998). In both models, performance levels at or near 80% could be achieved with as few as nine units in approximately 100 msec, a result consistent with the low convergence at early visual processing stages.

Figure 7 about here

So far, we have focused on relatively coarse discriminations between large and small spots. Artificially generated spike trains were used to study whether synchronous oscillations could, in principle, support finer discriminations over a wider range of spot sizes (fig. 8). For model I, spot sizes with L-values ranging from1 to 16 (see methods) were compared pairwise, while for model II, spot sizes ranged from 1×1 to 16×16 ganglion cells. The size of the smaller spot in each pair is given by the intercept of each curve with the x-axis, whereas the size of the larger spot is given by the x-coordinate of each point. Moving left to right along each curve, the size of the smaller spot is held

constant while the size of the larger spot increases. The analysis window for all pairwise size discriminations was 200 msec. Both models suggest that the size information encoded by synchronous oscillations, at least on behaviorally relevant time scales, is rather coarse, allowing discrimination only between small and larger spots. This coarse encoding was evident whether the multi-unit spike trains used in the discrimination analysis contained 4 cells (left column) or 9 cells (right column). Only poor discrimination (<75%) was possible between pairs of spots whose diameters both exceeded approximately 4-6 Poisson units or ganglion cells. Discrimination between large spots was slightly better in the Poisson model, for which the spectral amplitude and width depended linearly on spot size (see methods), whereas non-linear saturation effects were more significant in the network model. Even in the Poisson model, however, the firing probability could not drop below zero, which introduced a form of non-linear saturation. Such non-linear effects likely accounted for the decline in the performance mediated by both models as the second spot became very large. While these results must await experimental confirmation, they nonetheless suggest that the size information conveyed by synchronous oscillations is rather coarse, permitting discrimination only between small and large spots.

Figure 8 about here

The fact that retinal oscillations increase strongly with stimulus size suggests the total gamma-band activity, averaged across the entire retina, will exhibit a sharp dependence on the average pixel density. At low pixel densities, the total gamma-band activity should be negligible, as the image consists of only a few small, isolated spots. At very high pixel densities, the total gamma-band activity should be large, as an

appreciable fraction of the image will consist of a single contiguous spot. Interestingly, we expect that the total gamma-band activity will not depend smoothly on the average pixel density, but rather will exhibit a sharp transition, jumping discontinuously from small to large values, at a critical pixel density corresponding to a percolation transition.

To provide an intuitive feeling for a percolation transition, imagine a random binary image as a shallow, rock-strewn pond where the black pixels represent water and the white pixels represent rocks. Assuming that the rocks are distributed randomly, where p denotes the average density, we want to know whether it is possible, for a given value of p, to walk from one side of the pond to the other without getting wet. For small p, there will be only a few isolated rocks and an entirely dry passage is clearly impossible. As p increases, more and more rocks appear, forming small, isolated islands or clusters. Remarkably, at a very specific value of p, called the percolation threshold, the individual clusters tend to be joined together such that it suddenly becomes possible to walk all the way across the lake. The abrupt appearance of large connected clusters is the hallmark of a percolation transition, originally studied as a model of flow through porous media and which has a long history in statistical physics (e.g. Grimmett, 1999).

We used the retinal network model to investigate whether the total gamma-band activity, averaged across the entire image, exhibits a sharp transition around the percolation threshold. Random binary images were constructed by independently turning each pixel white, against a black background, with a probability p between zero and one (fig 9a). Above the percolation threshold, $p>p_c$, $p_c \approx 0.6$, there exists connected clusters that span the entire image. We computed trial-averaged power spectra from massed spike trains consisting of all 128×128 ganglion cells in the model (the size of the model was

increased for these experiments to reduce finite-size effects). As expected, there was a sharp increase in mean spectral amplitudes within the upper gamma-band as the pixel density crossed the percolation threshold, implying that synchronized retinal oscillations encode the presence of large features regardless of shape (fig. 9b). Both the average gamma-band activity (fig. 9c, solid line) and the average cluster size (see for example Hoshen and Kopelman, 1976) (fig. 9c, dotted dark gray line) exhibited sharp transitions around the percolation threshold, whereas the average luminance was simply proportional to the mean pixel density (fig. 9c, dashed light gray line). Retinal processing is thus predicted to convert a phase transition in the spatial properties of the input image into a transition in the temporal patterning of the output spike trains.

Figure 9 about here

A sharp increase in gamma-band activity was still evident near the percolation threshold even when the mean luminance was kept constant by reducing the intensity of each pixel in proportion to the mean pixel density (not shown), although the transition in this case was smoother due to the strong dependence of retinal oscillations on stimulus contrast (Kenyon et al., 2003; Kenyon et al., accepted for publication; Neuenschwander et al., 1999). The network model also lacked known physiological mechanisms for luminance adaptation that would likely have reduced the importance of the mean image intensity. The transition in gamma-band activity was not as sharp as the transition in mean cluster size, presumably because lateral processing in the retina tends to lump together large, nearly connected clusters (Kenyon et al., in press). For this reason, it was not necessary to consider different types of percolation (i.e. site or bond percolation) relating to the different ways connected clusters can be defined.

Discussion

We have examined the role of synchronous oscillations between retinal ganglion cells in the encoding of visual information. We find that high frequency oscillations in small groups of neighboring ganglion cells allow single-trial discrimination between large and small spots on biologically reasonable time scales. Single trial discrimination with low spatial convergence is essential if synchronous oscillations between retinal neurons are to convey stimulus attributes that are to be extracted early in the visual processing hierarchy. Similar single-trial discrimination was observed in two very different theoretical models whose multi-trial correlations were both consistent with experimental data, suggesting that the information conveyed by coherent oscillations is not strongly dependent on the underlying physiological mechanisms giving rise to the oscillations. In particular, similar single trial discrimination was obtained regardless of whether temporal correlations were produced by periodically modulated common input or through the dynamics of synaptically mediated feedback loops.

Ganglion cell firing rates are strongly modulated by local stimulus properties, such as contrast. The average firing rate among a small group of neighboring ganglion cells cannot therefore convey information about global stimulus properties without being confounded by local attributes. High frequency oscillations, on the other hand, are not evoked by small spots, even at high contrast. Thus, oscillations among a small group of neighboring ganglion cells can reliably signal that the stimulus to which they are responding is part of a large feature extending beyond their central receptive fields. High frequency oscillations therefore allow global topological information to be encoded by local spiking activity.

Encoding the presence of a large visual feature early in the visual system has potential behavioral advantages. For many organisms large objects represent possible danger and early detection allows for an early evasive response. It is possible, for instance, that high frequency oscillations between dimming detectors in the frog retina contribute to just such an early warning system (Ishikane et al., 1999). Likewise, in weakly electric fish, synchronized oscillations allow discrimination between prey and communication signals (Doiron et al., 2003). The question of how oscillatory information might be decoded by downstream networks is beyond the scope of the present study. However, we may imagine that corresponding high frequency resonances, of cellular, synaptic and/or network origin might cause neural circuits in the LGN or primary visual cortex to respond preferentially to oscillatory input.

The role of correlations in the neural code is controversial. Using an information theoretic analysis it has recently been argued that correlations between retinal ganglion cells encode relatively small amounts of additional information about natural visual scenes compared to their independent firing rates (Nirenberg et al., 2001). These results are not inconsistent with our work as their analysis used a 10 msec time window that could not have resolved oscillations in the upper gamma band. In addition, the role of synchrony as a possible binding mechanism has sparked intense debate (Shadlen and Movshon, 1999; Singer and Gray, 1995). Our results suggest that oscillations may encode global stimulus information regardless of their possible contribution to feature binding. Finally, it has been observed that any correlation between neural spike-trains is detrimental to rate-coded signals (Shadlen and Newsome, 1998). Correlations, however, are ubiquitous throughout the brain and our results demonstrate that, at least in some

cases, temporal correlations convey potentially useful information not present in the local mean firing rates.

It is possible to imagine additional physiological and behavioral experiments that might be conducted to further test the role of synchronous oscillations in the encoding of global topological properties. Indeed, our prediction of a sharp transition in total gamma-band activity as the pixel density crosses a percolation threshold suggests a novel and indirect method for examining the encoding role of synchronous oscillations. Presumably, local field potentials, or even clinical measures, such as the electroretinogram (ERG), could be used to assess gamma-band activity across large sections of the retina and there may be perceptual correlates of the predicted transition as well.

Appendix

Retinal Model

Artificial spike trains with realistic spatiotemporal correlations were generated by a retinal feedback circuit (fig 1), organized as a 32×32 array with wrap-around boundary conditions containing 5 distinct cell types: bipolar cells, small amacrine cells, large amacrine cells, axon-bearing amacrine cells, and ganglion cells. All cell types were modeled as single compartment, RC circuit elements obeying a first order differential equation that can be written efficiently in terms of matrix multiplications:

$$\dot{\vec{V}}^{(k)} = -\frac{1}{\tau^{(k)}}\left(\vec{V}^{(k)} - \sum_{k'} \vec{W}^{(k,k')} \cdot f^{(k,k')}\left(\vec{V}^{(k')}\right) \cdot \vec{W}^{(k,k')^T} - b^{(k)} - \vec{L}^{(k)} \right), \qquad \text{(A1)}$$

where $\vec{V}^{(k)}$ is a 2-D array denoting the normalized membrane potentials of all cells of type k, $(1 \le k \le 5)$, $\tau^{(k)}$ is the time constant, $b^{(k)}$ is a bias current for setting the resting potential, $\vec{L}^{(k)}$ is an external input intended to represent light stimulation, $\vec{W}^{(k,k')}$ gives the connection strengths between presynaptic $\{k'\}$ and postsynaptic $\{k\}$ cell types as a function of their separation along one direction, defined here as 'vertical', $\vec{W}^{(k,k')^T}$ gives the same information as a function of separation along the perpendicular direction, defined here as 'horizontal', and the functions $f^{(k,k')}$ give the associated input-output relations for the indicated pre- and post-synaptic cell types, detailed below. All membrane potentials were subject to a lower cutoff, equal to –1.5. The output of the axon-mediated inhibition was delayed by 2 msec, except for the axonal connections onto the axon-bearing amacrine cells, which was delayed for 1 msec. All other synaptic interactions were delayed by 1 msec. All equations were integrated in Matlab® using a direct Euler method with an integration time step of 1 msec.

The input-output function for gap junctions was given by the identity:

$$f^{(k,k')}\left(\vec{V}^{(k')}\right) = \vec{V}^{(k')}, \qquad \text{(A2)}$$

where the dependence on the presynaptic potential has been absorbed into the definition of $\tau^{(k)}$. This is possible because both the decay term in equation A1 and the omitted dependence on the presynaptic potential in equation A2 depend linearly on $\vec{V}^{(k)}$, allowing the coefficients to be combined. The input-output function for graded synapses was constructed by comparing, on each time step, a random number with a Fermi-function:

$$f^{(k,k')}\left(\vec{V}^{(k')}\right) = \theta\left(\left[\frac{1}{1+\exp\left(-\alpha\vec{V}^{(k')}\right)}\right] - r\right), \tag{A3}$$

where α sets the gain (equal to 4 for all non-spiking synapses), r is a uniform random deviate equally likely to take any real value between 0 and 1, and θ is a step function, $\theta(x) = 1,\ x \geq 0\ ;\ \theta(x) = 0,\ x < 0$.

Lastly, the input-output relation used for conventional threshold synapses was:

$$f^{(k,k')}\left(\vec{V}^{(k')}\right) = \theta\left(\vec{V}^{(k')}\right). \tag{A4}$$

A modified integrate-and-fire mechanism was used to model spike generation. A positive pulse (amplitude = 10.0) was delivered to the cell on the time step after the membrane potential crossed threshold, followed by a negative pulse (amplitude = -10.0) on the subsequent time step. This resulted in a 1 msec action potential that also produced impulse responses in electrically coupled cells, an important element of the circuit dynamics. The bias current, *b*, was incremented by -0.5 following each spike, and then decayed back to the resting value with the time constant of the cell, adding to the relative refractory period. There was in addition an absolute refractory period of 1 msec.

Along both the horizontal and vertical directions, synaptic strengths fell off as Gaussian functions of the distance between the pre- and post-synaptic cells. For a given horizontal separation, the horizontal weight factor was determined by a Gaussian function of the following form:

$$W_{i^{(k)},j^{(k')}}^{(k,k')} = \alpha\sqrt{W^{(k,k')}}\exp\left[-\frac{\left\|i^{(k)} - j^{(k')}\right\|^2}{2\sigma^2}\right] \quad \text{(A5)}$$

where $W_{i^{(k)},j^{(k')}}^{(k,k')}$ is the horizontal weight factor from presynaptic cells of type k' located in the j^{th} column to the postsynaptic cells of type k located in the i^{th} column, α is a normalization factor, determined numerically, which ensured that the total synaptic input integrated over all presynaptic cells of type k' to every postsynaptic cell of type k equaled $W^{(k,k')}$, σ is the Gaussian radius of the interaction, and the quantity $\left\|i^{(k)} - j^{(k')}\right\|$ denotes the horizontal distance between the pre- and post-synaptic cells, taking into account the wrap around boundary conditions employed to mitigate edge effects. An analogous weight factor describes the dependence on vertical separation. Equation A5 was augmented by a cutoff condition that prevented synaptic interactions beyond a specified distance, determined by the radius of influence of the presynaptic outputs and the postsynaptic inputs, roughly corresponding to the axonal and dendritic fields, respectively. A synaptic connection was only possible if the output radius of the presynaptic cell overlapped the input radius of the postsynaptic cell. Except for axonal connections, the input and output radii were the same for all cell types. For the large amacrine cells and the ganglion cells, the radius of influence extended out to the centers of the nearest neighboring cells of the same type, producing a coverage factor greater than one (Vaney, 1990). The radii of the bipolar, small, and axon-bearing amacrine cells (non-axonal connections only) extended only halfway to the nearest cell of the same type, giving a coverage factor of one (Cohen and Sterling, 1990). The external input was multiplied by a gain factor of 3. Values for model parameters are listed in tables 1 and 2. A more complete description of the model is presented elsewhere (Kenyon et al., 2003).

Table 1: Cellular parameters.

	τ	b	n×n	d	σ
BP	10.0	-0.0	64×64	0.25	0.25
SA	25.0	-0.5	64×64	0.25	0.25
LA	20.0	-0.25	32×32	1.0	0.5
PA	5.0	-0.025	64×64	0.25/9.0[a]	0.25/3.0[a]
GC	5.0	-0.025	32×32	1.0	0.5

Explanation of symbols: τ: time constant (msec); b: bias; n×n: array size; d: cutoff radius, σ: Gaussian radius (see eq. 5). [a]Inner radius/outer radius.

Table 2: Synaptic weights.

	BP	SA	LA	PA	GC
BP	*	-0.375[b]	-3.0b	-3.0b/-15.0c	*
SA	3.0[b]	*	-3.0[b]	0.0b/-15.0[c]	*
LA	3.0[b]	*	0.25[a]	-3.0a/-15.0[c]	*
PA	0.75[b]	-0.75b	0.25[a]	0.25a/-45.0[c]	0.25[a,d]
GC	9.0[b]	-4.5b	-4.5[b]	0.25[a]/-270.0[c]	*

Each term represents the total integrated weight from all synapses arising from the corresponding presynaptic type (columns) to each cell of the corresponding postsynaptic type (rows), (the quantity $W^{(k,k')}$ in eq. 5). Asterisk (*) indicates absence of corresponding connection. Synapse type indicated by superscript: [a]gap junction, [b]non-spiking synapse, [c]graded synapse. [d]Maximum coupling efficiency (ratio of post- to pre-synaptic depolarization) for this gap junction synapse: DC=11.3%, Action Potential=2.7%.

Figure Captions

Figure 1. Schematic of the retinal feedback circuit. Only those elements directly responsible for synchronous oscillations are depicted. A combination of local excitation via gap junctions (resistors) and long-range inhibition via axon-bearing amacrine cells (gray dotted lines and filled black circles) produced physiologically realistic size-dependent oscillations. The entire model contained a 32×32 array of ganglion cells driven by a 64×64 array of bipolar cells, with inhibitory feedback arising from a 64×64 array of axon-bearing amacrine cells, only a few of which are shown. Light stimuli were implemented by injecting currents directly into bipolar cells.

Figure 2. Multi-unit auto- and cross-correlograms reveal size-dependent high frequency synchronous oscillations. a) Auto-correlograms were computed from multi-unit spike trains recorded from cat retina, with the electrode located at the center of square spots of increasing size (data replotted from (Neuenschwander et al., 1999)). Correlations expressed as a fraction of the expected level due to chance. Bottom Trace: spot size = 0.7º×0.7º. Middle Trace: spot size = 6.3º×6.3º. Top Trace: spot size = 9.8º×9.8º. Individual traces are offset for clarity. (b) Multi-unit cross-correlograms of artificial spike trains generated by a rate-modulated Poisson process containing 4 identical units, each with a mean firing rate of 50 Hz. c) Multi-unit cross-correlograms produced by an integrate-and-fire feedback circuit consistent with retinal anatomy. The retinal model was stimulated by stationary spots ($\log_2$ intensity = -2) covering 1×1, 4×4 and 6×6 ganglion cells. Multi-unit spike-trains were recorded from a fixed 2×2 array of ganglion

cells located at the center of each spot. Poisson distributed spikes were added which increased the mean firing rate by 20% regardless of spot size, a procedure that resulted in a better fit to the physiological data.

Figure 3. Frequency spectra of natural and artificial multi-unit spike trains as a function of spot size. Same data and general organization as in figure 2. a) Frequency spectra, normalized by the total number of spikes, computed from multi-unit spike trains recorded from the cat retina in response to different sized spots. b) Normalized frequency spectra of multi-unit spike trains generated by a rate-modulated Poisson process. c) Normalized frequency spectra of multi-unit spike trains produced by the retinal feedback circuit. In all panels, a prominent peak in the upper gamma band, 65-100 Hz (dotted vertical lines), becomes evident as spot size increases.

Figure 4. Single trial frequency spectra of natural and artificial multi-unit spike trains as a function of spot size. Same data and general organization as in figure 3. Individual trials are drawn as lines with different shading. Although considerably noisier than their trial-averaged counterparts, an increase in spectral amplitudes in the upper gamma band is still clearly evident as a function of increasing stimulus size, suggesting that synchronous oscillations encode topological information that is available on behaviorally relevant time scales.

Figure 5. Distribution of single-trial gamma-band activity as a function of stimulus size. Same data as in figures 2 and 3. Gamma-band activity, defined as the mean (normalized)

Fourier amplitude between 65 and 100 Hz relative to the mean amplitude between 220 and 500 Hz, was sorted into uniform bins (arbitrary units). a) Distribution of single-trial gamma-band activity recorded from cat retinal ganglion cells. The distribution of gamma-band activity elicited by a small spot (solid line) is largely distinguishable from that elicited by both intermediate (dashed line) and large spots (solid line). b) Distributions of gamma-band activity for the rate-modulated Poisson process. c) Analogous distributions for the retinal network. In both sets of artificially generated spike trains, the distributions associated with small vs. intermediate and small vs. large spot sizes are clearly distinct.

Figure 6. Theoretically optimal performance on a size discrimination task. Same data as in figure 5. A Bayes discriminator was used to categorize relative spot size, as either "smaller" or "larger", based on the single-trial gamma-band activity. Two pairs of spot sizes were employed, either small vs. medium or small vs. large. The sizes of the small, medium and large spots were as indicated in previous figures. The height of each bar gives the maximum percentage of trials that could be classified correctly, assuming both sizes were equally likely *a priori*, based on the distributions shown in figure 5. All three data sets indicate that synchronous oscillations within a small group of ganglion cells are capable of supporting reasonable discrimination between small and large spots on behaviorally relevant time scales.

Figure 7. Size-discrimination improves as multi-unit spike train segments become longer and/or include more cells. A Bayes discriminator was used to classify spots as being

either “smaller” or “larger” based on the single-trial gamma-band activity. a) Maximum percentage of correctly classified trials plotted as a function of the analysis window size for multi-unit spike trains recorded from cat retinal ganglion cells. Performance levels ranged between approximately 70 to 75% as the length of the multi-unit spike train segment increased from 50 to 300 msec. b) Percent of correctly classified trials using multi-unit spike trains generated by a rate-modulated Poisson process. Performance improved with longer analysis windows and as more units were included in the spike record. (Explanation of symbols: ○=1, □=4, ◊=9, ×=16 units). c) Percentage of correctly classified trials using multi-unit spike trains generated by the retinal network. (Explanation of symbols: ○=1×1, □=2×2, ◊=3×3, ×=4×4 model ganglion cells). Performance again improved as the analysis window size and the number of neighboring ganglion cells included in the analysis was increased.

Figure 8. Synchronous oscillations do not support fine discrimination between spot sizes. Pairs of stimuli, consisting of two different sized spots, were classified as either “smaller” or “larger” based on the single-trial gamma-band activity. The maximum percent correct, estimated via Bayesian discrimination, is plotted for all pairwise combinations. Each point represents a different discrimination task. The size of the smaller spot is given by the x-intercept of the containing curve and the size of the larger spot by the x-coordinate. a,c) Rate-modulated Poisson process consisting of a) 4 units and c) 9 units. Spot sizes specified as L-values (see methods) ranging from 1 to 16. Sizes corresponding to the labels “small”, “med” and “large” used in previous figures are indicated separately. b,d) Multi-unit spike trains from the retinal network model consisting of b) 2×2 and d) 3×3

neighboring ganglion cells. Performance on the pairwise size discrimination task was poor between spots containing more than 4-9 cells, suggesting that synchronous oscillations primarily allow only coarse discriminations between small and large spots.

Figure 9. Binary images produce a sharp transition in gamma-band activity at the percolation threshold. a) Random binary images with varying pixel density, p, indicated above each image. Above the percolation threshold ($p \approx 0.6$), connected clusters of white pixels span the image. For clarity, the total number of pixels was reduced in these illustrations from 256×256 to 64×64. b) Trial-averaged multi-unit frequency spectra computed from massed spike trains combining all 128×128 model ganglion cells. Above the percolation threshold, a large peak appears in the upper gamma band. c) Gamma-band activity (solid black line) exhibits a sharp transition around the percolation threshold, consistent with the phase transition in the mean size of connected clusters (dotted dark gray line). Total luminance (dashed light gray line), which grows linearly with pixel density, cannot account for the sharp transition in gamma-band activity.

References

Alonso J.M., Usrey W.M., Reid R.C. (2001). Rules of connectivity between geniculate cells and simple cells in cat primary visual cortex. J Neurosci, 21, 4002-4015.

Ariel M., Daw N.W., Rader R.K. (1983). Rhythmicity in rabbit retinal ganglion cell responses. Vision Research, 23, 1485-1493.

Cohen E., Sterling P. (1990). Convergence and divergence of cones onto bipolar cells in the central area of cat retina. Philosophical Transactions of the Royal Society of London - Series B: Biological Sciences, 330, 323-328.

Dacey D.M., Brace S. (1992). A coupled network for parasol but not midget ganglion cells in the primate retina. Visual Neuroscience, 9, 279-290.

De Carli F., Narici L., Canovaro P., Carozzo S., Agazzi E., Sannita W.G. (2001). Stimulus- and frequency-specific oscillatory mass responses to visual stimulation in man. Clin Electroencephalogr, 32, 145-151.

Doiron B., Chacron M.J., Maler L., Longtin A., Bastian J. (2003). Inhibitory feedback required for network oscillatory responses to communication but not prey stimuli. Nature, 421, 539-543.

Duda R.O., Hart P.E., Stork D.G. (2001). Pattern Classification. NY: Wiley.

Freed M.A. (2000). Rate of quantal excitation to a retinal ganglion cell evoked by sensory input. J Neurophysiol, 83, 2956-2966.

Frishman L.J., Saszik S., Harwerth R.S., Viswanathan S., Li Y., Smith E.L., 3rd, Robson J.G., Barnes G. (2000). Effects of experimental glaucoma in macaques on the

multifocal ERG. Multifocal ERG in laser-induced glaucoma. Doc Ophthalmol, 100, 231-251.

Grimmett G. (1999). Percolation. Berlin ; New York: Springer.

Hoshen J., Kopelman R. (1976). Percolation and cluster distribution. I. cluster multiple labeling technique and critical concentration algorithm. Phys Rev B, 14, 3438-3445.

Ishikane H., Kawana A., Tachibana M. (1999). Short- and long-range synchronous activities in dimming detectors of the frog retina. Vis Neurosci, 16, 1001-1014.

Jacoby R., Stafford D., Kouyama N., Marshak D. (1996). Synaptic inputs to ON parasol ganglion cells in the primate retina. Journal of Neuroscience, 16, 8041-8056.

Kenyon G.T., Moore B., Jeffs J., Denning K.S., Stephens G.J., Travis B.J., George J.S., Theiler J., Marshak D.W. (2003). A model of high-frequency oscillatory potentials in retinal ganglion cells. Vis Neurosci, 20, 465-480.

Kenyon G.T., Theiler J., Travis B.J., George J.S., Marshak D.W. (accepted for publication). Correlated firing improves stimulus discrimination. Neural Computation.

Kenyon G.T., Travis B.J., George J.S., Theiler J., Marshak D.W. (in press). Stimulus-specific oscillations in a retinal model. IEEE Transactions on Neural Networks: Special Issue on Temporal Coding for Neural Information Processing.

Laufer M., Verzeano M. (1967). Periodic activity in the visual system of the cat. Vision Res, 7, 215-229.

Martinez-Conde S., Macknik S.L., Hubel D.H. (2004). The role of fixational eye movements in visual perception. Nat Rev Neurosci, 5, 229-240.

Neuenschwander S., Castelo-Branco M., Singer W. (1999). Synchronous oscillations in the cat retina. Vision Res, 39, 2485-2497.

Neuenschwander S., Singer W. (1996). Long-range synchronization of oscillatory light responses in the cat retina and lateral geniculate nucleus. Nature, 379, 728-732.

Nirenberg S., Carcieri S.M., Jacobs A.L., Latham P.E. (2001). Retinal ganglion cells act largely as independent encoders. Nature, 411, 698-701.

O'Brien B.J., Isayama T., Richardson R., Berson D.M. (2002). Intrinsic physiological properties of cat retinal ganglion cells. J Physiol, 538, 787-802.

Shadlen M.N., Movshon J.A. (1999). Synchrony unbound: a critical evaluation of the temporal binding hypothesis. Neuron, 24, 67-77, 111-125.

Shadlen M.N., Newsome W.T. (1998). The variable discharge of cortical neurons: implications for connectivity, computation, and information coding. J Neurosci, 18, 3870-3896.

Singer W., Gray C.M. (1995). Visual feature integration and the temporal correlation hypothesis. Annual Review of Neuroscience, 18, 555-586.

Steinberg R.H. (1966). Oscillatory activity in the optic tract of cat and light adaptation. J Neurophysiol, 29, 139-156.

Usrey W.M., Reppas J.B., Reid R.C. (1999). Specificity and strength of retinogeniculate connections. J Neurophysiol, 82, 3527-3540.

Vaney D.I. (1990) The mosaic of amacrine cells in the mammalian retina. In: Progress in Retinal Research (Osborne N.N., Chader G.J., eds), pp 49-100. Oxford: Pergamon Press.

Vaney D.I. (1994). Patterns of neuronal coupling in the retina. Progress in Retinal and Eye Research, 13, 301-355.

Wachtmeister L. (1998). Oscillatory potentials in the retina: what do they reveal. Prog Retin Eye Res, 17, 485-521.

Wachtmeister L., Dowling J.E. (1978). The oscillatory potentials of the mudpuppy retina. Invest Ophthalmol Vis Sci, 17, 1176-1188.

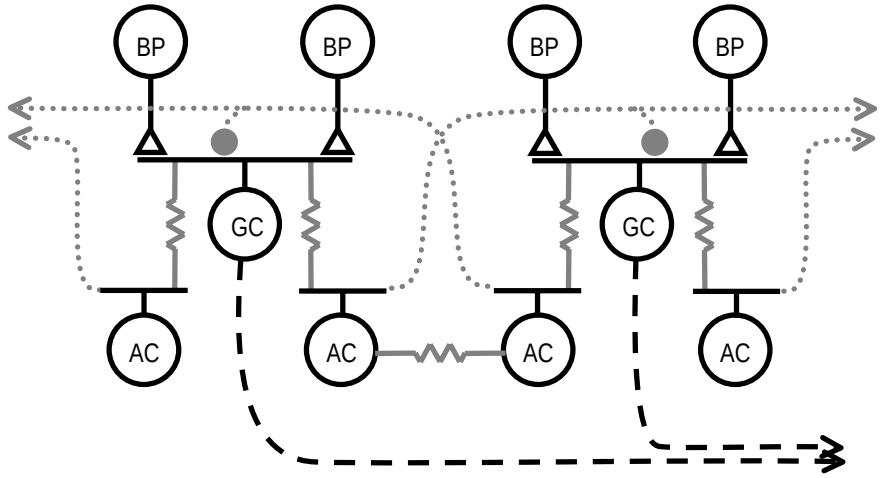

Fig 1, Kenyon, NC 2868

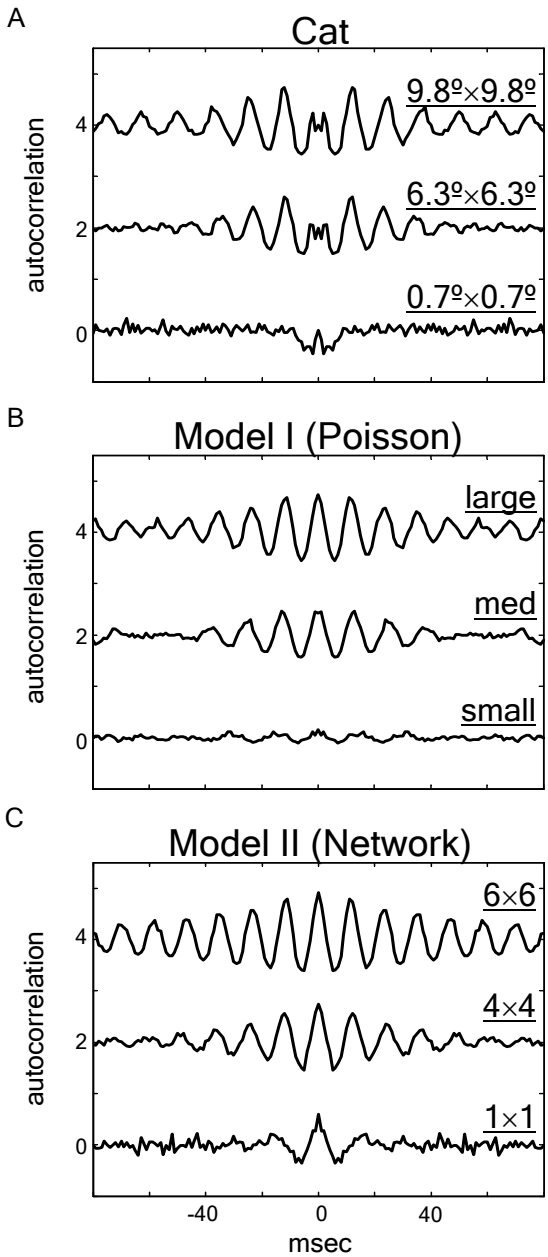

Fig 2, Kenyon, NC 2868

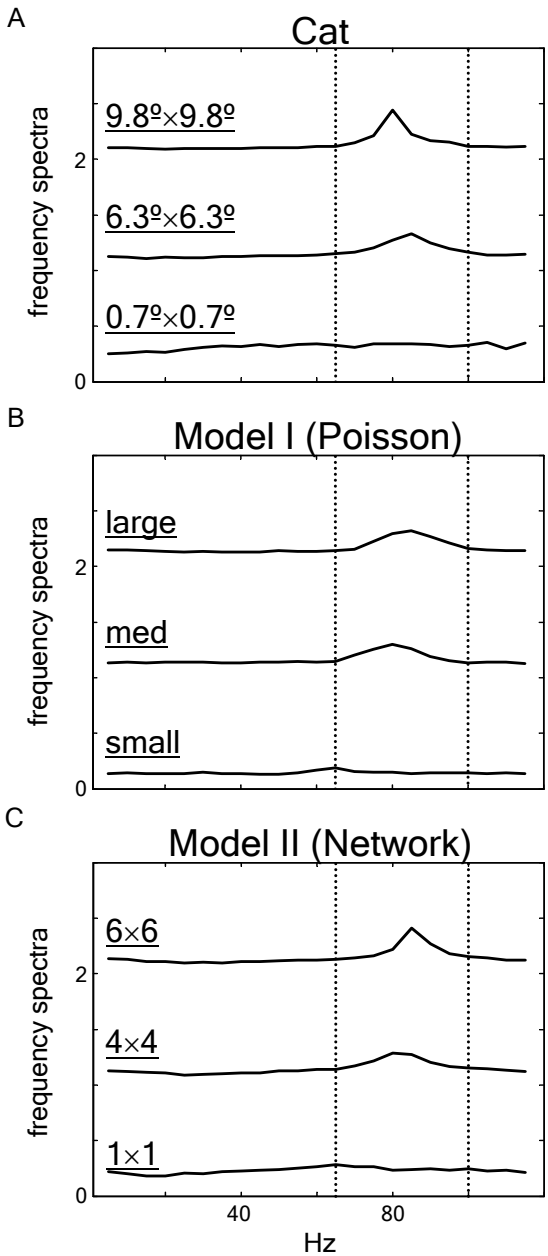

Fig 3, Kenyon, NC 2868

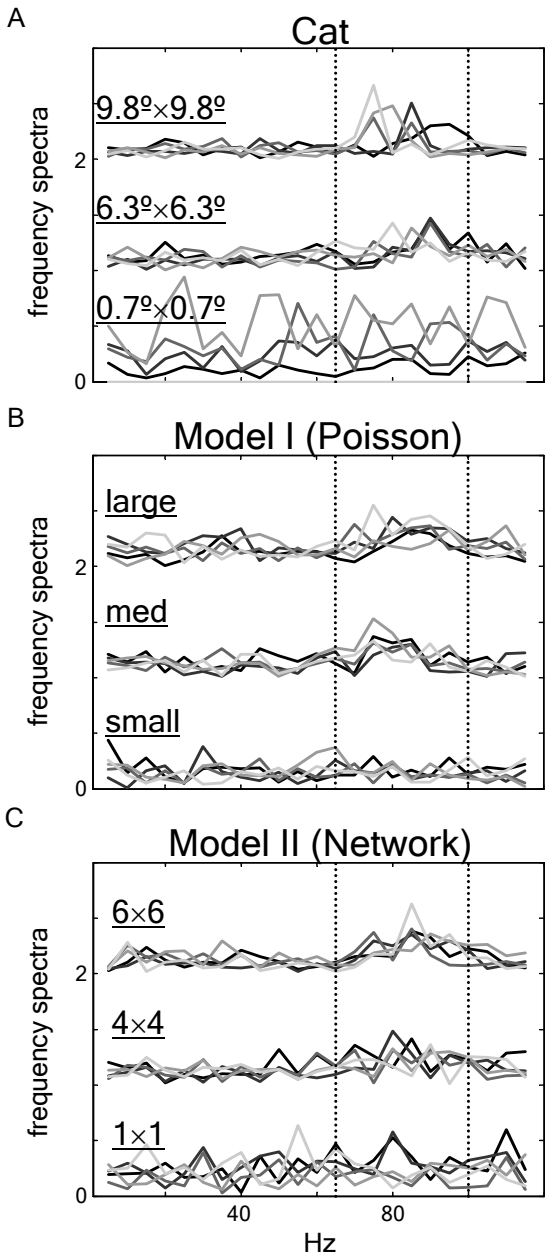

Fig 4, Kenyon, NC 2868

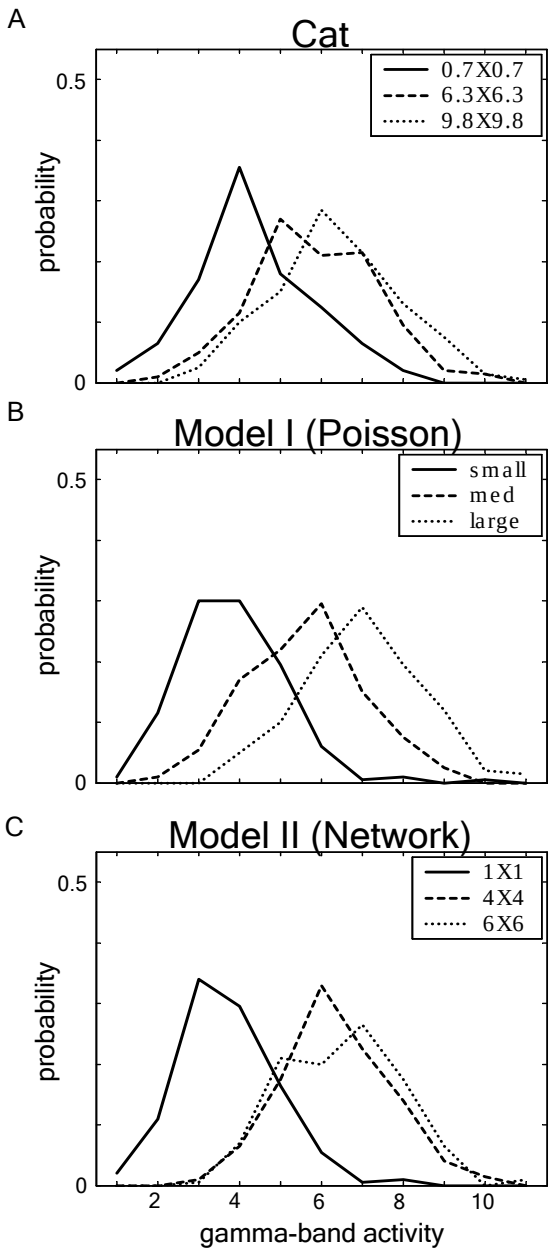

Fig 5, Kenyon, NC 2868

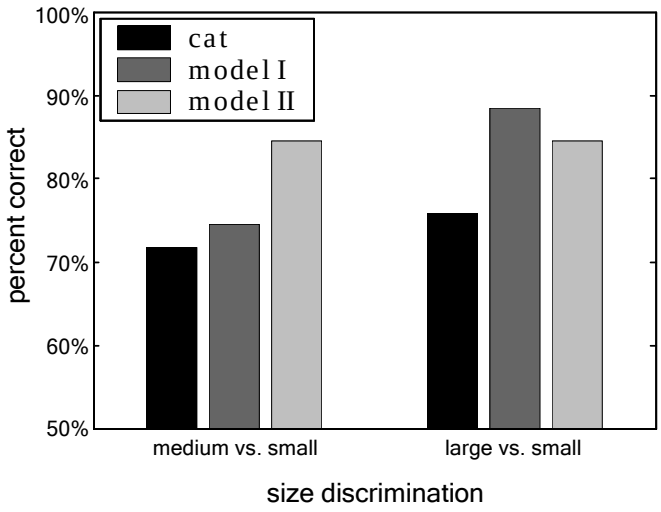

Fig 6, Kenyon, NC 2868

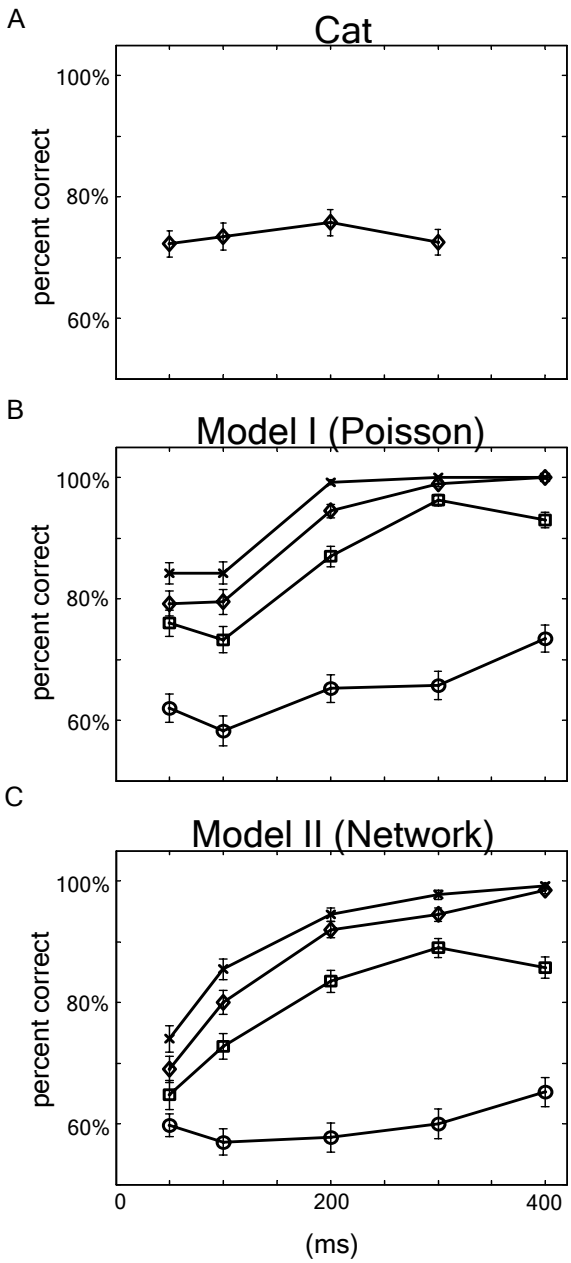

Fig 7, Kenyon, NC 2868

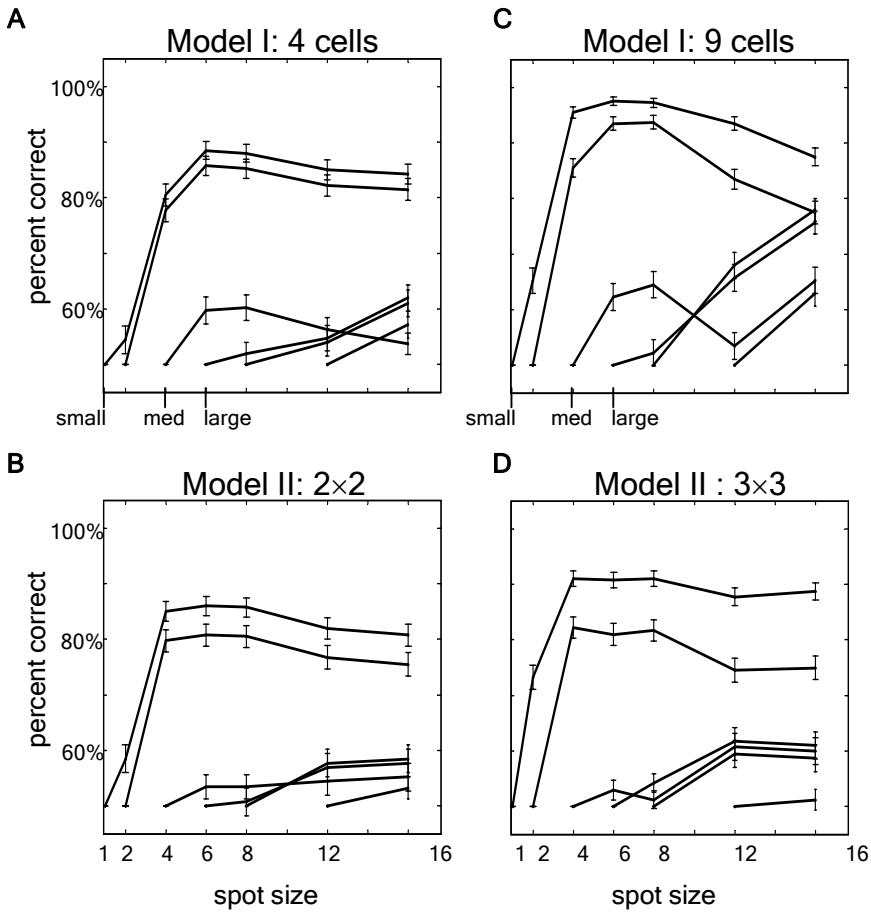

Fig 8, Kenyon, NC 2868

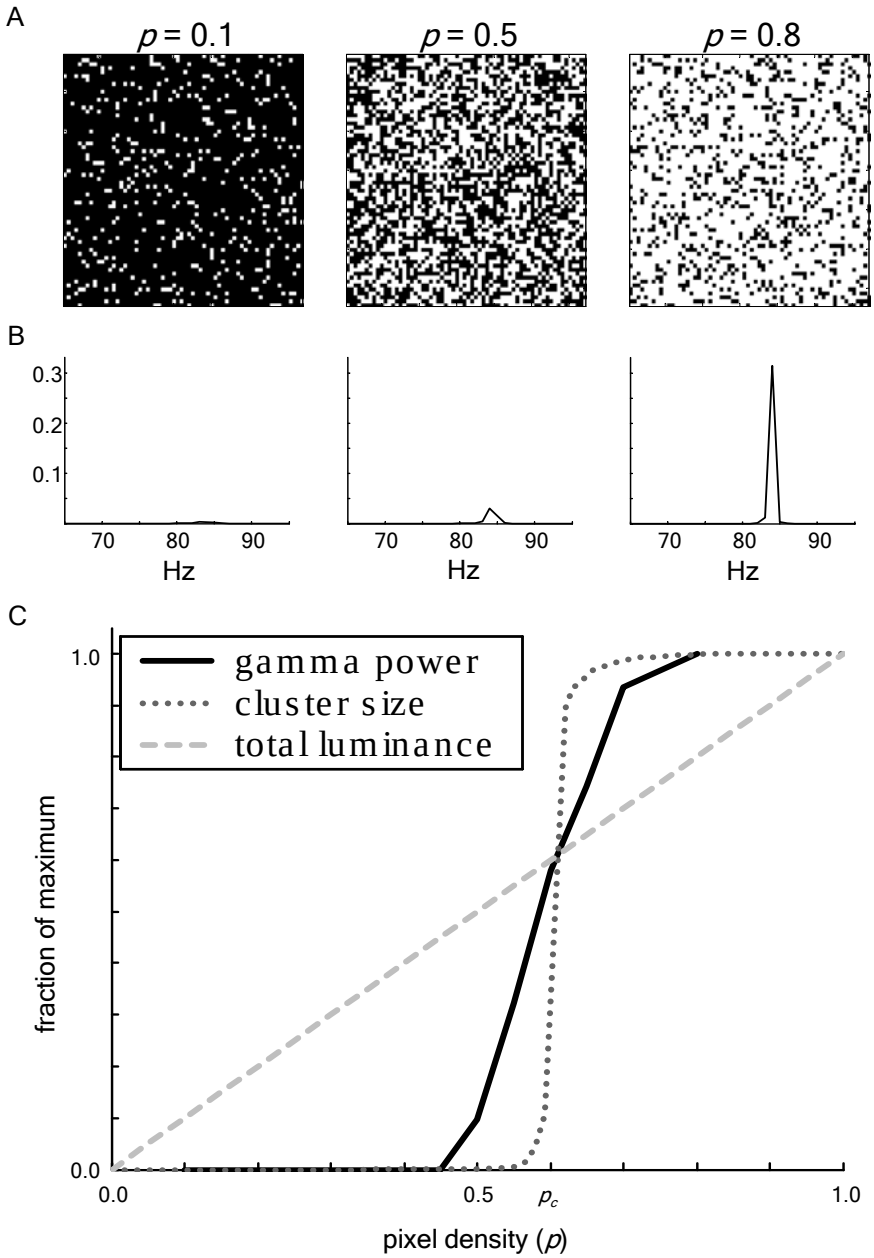

Fig 9, Kenyon, NC 2868